Chapter 8. Rigid transformations

We are about to start drawing figures in 3D. There are no built-in routines for this purpose in PostScript, and we shall have to start more or less from scratch in extending the language to do the job we want done. Figures in 3D are a great deal more complicated than ones in 2D, and there are a number of new mathematical ideas to be introduced.

The biggest problem in 3D is you can't represent an object exactly, since in drawing it you must collapse it to 2D. Nonetheless this aim is not entirely unreasonable, because of course our eyes render the world on the back of the retina, itself a two-dimensional surface. But what we need to do to make the illusion work is to move objects around as we look at them, or at least allow for this possibility; and introduce shadows and other responses to light to create a perception of depth.

In the next chapter we shall describe in detail how we arrange viewing things in space. In this one we shall describe how things move around without distortion. It is useful to discuss this in dimensions one and two as well as three.

1. Rigid transformations

If we move an object around normally, it will in some sense remain rigid, and will not distort. Here is the technical way we formulate rigidity: Suppose we move an object from one position to another. In this process, any point P of the object will be moved to another point P_*. We shall say that the points of the object are **transformed** into other points. A transformation is said to be **rigid** if it preserves relative distances—that is to say, if P and Q are transformed to P_* and Q_* then the distance from P to Q is the same as that from P_* to Q_*.

We shall make an extra assumption about rigid transformations. It happens that it is a redundant assumption, since it can be proven that every rigid transformation satisfies this condition. We shall *not* prove it, however, because it would require a long digression, and instead just take it more or less for granted.

The condition is this: *All the rigid transformations we consider will be* **affine**. This means that if we have chosen a linear coordinate system in whatever set we are looking at (a line, a plane, or space). then the transformation $P \mapsto P_*$ is calculated in terms of coordinate vectors x and x_* according to the formula

$$x_* = Ax + v$$

where A is a matrix and v a vector. In 3D, for example, we require

$$\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = A \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} + \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} .$$

It turns out that all rigid transformations are in fact affine, but we shall not worry about that here. The matrix A is called the **linear component**, v the **translation component** of the transformation.

- *A rigid transformation preserves angles as well as distances.*

That is to say, if P, Q and R are three points transformed to P_*, Q_*, and R_*, then the angle θ between segments PQ and PR is the same as the angle θ_* between P_*Q_* and P_*R_*. This is because of the cosine law, which says that

$$\begin{aligned} \cos\theta &= \frac{\|QR\|^2 - \|PQ\|^2 - \|PR\|^2}{\|PQ\|\,\|PR\|} \\ &= \frac{\|Q_*R_*\|^2 - \|P_*Q_*\|^2 - \|P_*R_*\|^2}{\|P_*Q_*\|\,\|P_*R_*\|} \\ &= \cos\theta_* \ . \end{aligned}$$

A few other facts are more elementary:

- *The composition of rigid transformations is rigid.*
- *The inverse of a rigid transformations is rigid.*

In the second statement, it is implicit that a rigid transformation has an inverse. This is easy to demonstrate. An affine transformation will be rigid when its linear component is, since a translation will certainly not distort lengths. But if its linear component does not have an inverse, then it is **singular**, which means that it will collapse some line, at least, onto a point. Then it cannot preserve lengths, which is a contradiction.

In order to classify rigid transformations, which is what we shall now do, we must thus classify the linear ones.

Exercise 1.1. *The inverse of the transformation $x \mapsto Ax + v$ is also affine. What are its linear and translation components?*

2. Dot and cross products

A bit later we shall need to know some basic facts about vector algebra. In any number of dimensions we define the dot product of two vectors

$$u = (x_1, x_2 \ldots, x_n), \quad v = (y_1, y_2, \ldots, y_n)$$

to be

$$u \bullet v = x_1 y_1 + x_2 y_2 + \cdots x_n y_n \ .$$

The relation between dot products and geometry is expressed by the cosine rule for triangles, which asserts that if θ is the angle between u and v then

$$\cos\theta = \frac{u \bullet v}{\|u\| \, \|v\|} \ .$$

In particular u and v are perpendicular when $u \bullet v = 0$.

In $3D$ there is another kind of product. If

$$u = (x_1, x_2, x_3), \quad v = (y_1, y_2, y_3)$$

then their **cross product** $u \times v$ is the vector

$$(x_2 y_3 - y_2 x_3, x_3 y_1 - x_1 y_3, x_1 y_2 - x_2 y_1) \ .$$

This formula can be remembered if we write the vectors u and v in a 2×3 matrix

$$\begin{bmatrix} x_1 & x_2 & x_3 \\ y_1 & y_2 & y_3 \end{bmatrix}$$

and then for each column of this matrix calculate the determinant of the 2×2 matrix we get by crossing out in turn each of the columns. The only tricky part is that *with the middle coefficient we must reverse sign.* Thus

$$u \times v = \left(\begin{vmatrix} x_2 & x_3 \\ y_2 & y_3 \end{vmatrix}, - \begin{vmatrix} x_1 & x_3 \\ y_1 & y_3 \end{vmatrix}, \begin{vmatrix} x_1 & x_2 \\ y_1 & y_2 \end{vmatrix}, \right) \ .$$

Here

$$\begin{vmatrix} a & b \\ c & d \end{vmatrix} = ad - bc \ .$$

The geometrical significance of the cross product is contained in these rules:

- *The length of $w = u \times v$ is the area of the parallelogram spanned in space by u and v.*

- *It lies in the line perpendicular to the plane containing u and v and its direction is determined by the right hand rule.*

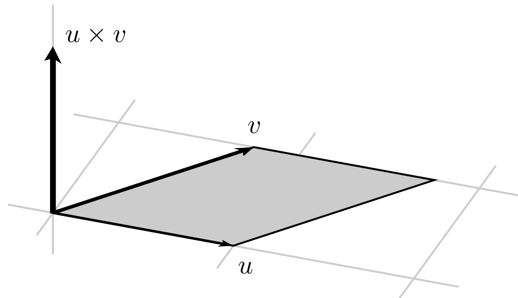

The cross product $u \times v$ will vanish only when u and v are multiples of one another.

In these notes, the main use of dot products and cross products will be in calculating **projections**.

(1) Suppose α to be any vector in space and u some other vector in space. The **projection of u along α** is the vector u_0 we get by projecting u perpendicularly onto the line through α.

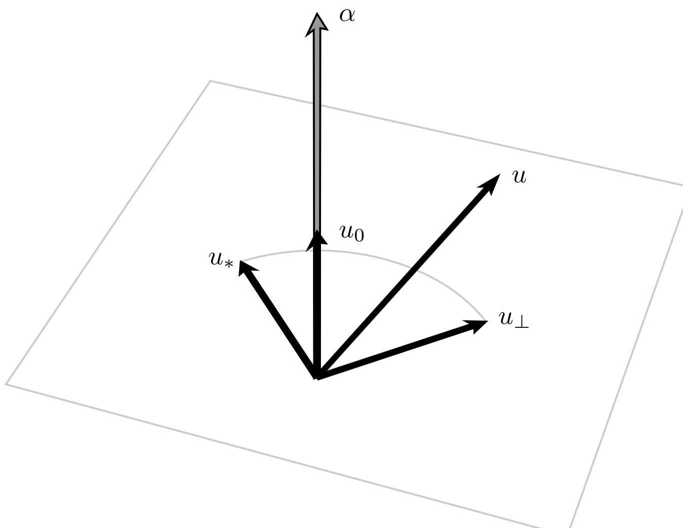

What is this projection? It must be a multiple of α. We can figure out what multiple by using trigonometry. We know three facts: (a) The angle θ between α and u is determined by the formula

$$\cos\theta = \frac{u \bullet \alpha}{\|\alpha\|\,\|u\|} \ .$$

(b) The length of the vector u_0 is $\|u\|\cos\theta$, and this is to be interpreted algebraically in the sense that if u_0 faces in the direction opposite to α this is negative. (c) Its direction is either the same or opposite to α. The vector $\alpha/\|\alpha\|$ is a vector of unit length pointing in the same direction as α. Therefore

$$u_0 = \|u\|\cos\theta\ \frac{\alpha}{\|\alpha\|} = \|u\|\ \frac{u \bullet \alpha}{\|\alpha\|\,\|u\|}\ \frac{\alpha}{\|\alpha\|} = \left(\frac{u \bullet \alpha}{\|\alpha\|^2}\right)\alpha = \left(\frac{u \bullet \alpha}{\alpha \bullet \alpha}\right)\alpha \ .$$

(2) Now let $u_\perp$ be the projection of u onto the plane Π through the origin perpendicular to α.

The vector u has the orthogonal decomposition

$$u = u_0 + u_\perp$$

and therefore we can calculate

$$u_\perp = u - u_0 \ .$$

(3) Finally, let u_* be the vector in Π we get by rotating u_* by 90° in Π, using the right hand rule to determine what direction of rotation is positive.

How do we calculate u_*? We want it to be perpendicular to both α and $u_\perp$, so it ought to be related to the cross product $\alpha \times u_\perp$. A little thought should convince you that in fact the direction of u_* will be the same as that of $\alpha \times u_\perp$, so that u_* will be a positive multiple of $\alpha \times u_\perp$. We want u_* to have the same length as $u_\perp$. Since α and $u_\perp$ are perpendicular to each other, the length of the cross product is equal to the product of the lengths of α and $u_\perp$, and we must divide by $\|\alpha\|$ to get a vector of length $\|u_\perp$. Therefore, all in all

$$u_* = \frac{\alpha}{\|\alpha\|} \times u_\perp \ .$$

Incidentally, in all of this discussion it is only the direction of α that plays a role. It is often useful to normalize α right at the beginning of these calculations, that is to say replace α by $\alpha/\|\alpha\|$.

Exercise 2.1. *Write PostScript programs to calculate dot products, cross products, u_0, $u_\perp$, u_*.*

3. The classification of linear rigid transformations

Let A be an $n \times n$ matrix.

$$A = \begin{bmatrix} a_{1,1} & a_{1,2} & \dots & a_{1,n} \\ a_{2,1} & a_{2,2} & \dots & a_{2,n} \\ & & \dots & \\ a_{n,1} & a_{n,2} & \dots & a_{n,n} \end{bmatrix} \ .$$

Let e_i be the column vector with exactly one non-zero coordinate equal to 1 in the i-th place. Then Ae_i is equal to the i-th column of A. The transformation corresponding to A takes the origin to itself, and the length of e_i is 1. Therefore the length of the i-th column of A is also 1.

The angle between e_i and e_j is 90° if $i \neq j$, and therefore the angle between the i-th and j-th columns of A is also 90°.

Since the square of the length of a vector u is equal to the dot product $u \bullet u$, and the angle between two vectors of length 1 is given by their dot product $u \bullet v$, this means that the columns u_i of A satisfy the relations

$$u_i \bullet u_j = \begin{cases} 1 & i = j \\ 0 & \text{otherwise} \end{cases}$$

Any matrix A satisfying these conditions is said to be **orthogonal**. The **transpose** ${}^tA$ of a matrix A is obtained by flipping A along its diagonal. In other words, the rows of ${}^tA$ are the columns of A, and vice-versa. By definition of the matrix product ${}^tA\,A$, its entries are the various dot products of the columns of A with the rows of ${}^tA$. Therefore a matrix A is orthogonal if and only if

$${}^tA\,A = I, \quad A^{-1} = {}^tA \ .$$

If A and B are two $n \times n$ matrices, then

$$\det(AB) = \det(A)\det(B) \ .$$

The determinant of A is the same as that of its transpose. If A is an orthogonal matrix, then

$$\det(I) = \det(A)\det({}^tA) = \det(A)^2$$

so that $\det(A) = \pm 1$. If $\det(A) = 1$, A is said to **preserve orientation**, otherwise **reverse orientation**. There is a serious qualitative difference between the two types. If we start with an object in one position and move it continuously, then the transformation describing its motion will be a continuous family of rigid transformations. The linear component at the beginning is the identity matrix, with determinant 1. Since the family varies continuously, the linear component can never change the sign of its determinant, and must therefore always be orientation preserving. A way to change orientation would be to reflect the object, as if in a mirror.

R Я

4. Orthogonal transformations in 2D

In 2D the classification of orthogonal transformations is very simple. First of all, we can rotate an object through some angle (possibly $0°$).

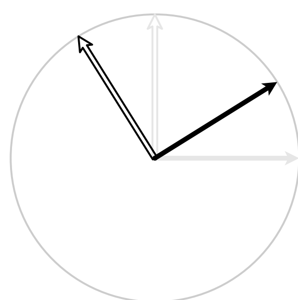

This preserves orientation. The matrix of this transformation is, as we saw much earlier,

$$\begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix} .$$

Second, we can reflect things in an arbitrary line.

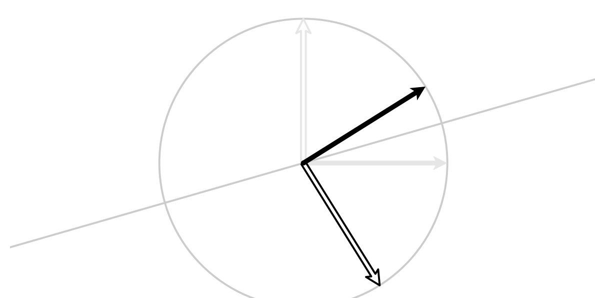

That is to say, given a line ℓ, we can transform points on ℓ into themselves, points in the line through the origin perpendicular to ℓ into their negatives. This reverses orientation.

Exercise 4.1. *If ℓ is the line at angle θ with respect to the positive x-axis, what is the matrix of this reflection?*

It turns out there are no more possibilities.

- *Every linear rigid transformation in 2D is either a rotation or a reflection.*

Let $e_1 = (1,0)$, $e_2 = (0,1)$, and let T be a linear rigid transformation. Since e_1 and e_2 both have length 1, both Te_1 and Te_1 also have length 1. All of these lie on the unit circle. Since the angle between e_1 and e_2 is $90°$, so is that between Te_1 and Te_2. There are two distinct possibilities, however. Either we rotate in the positive direction from Te_1 to Te_2, or in the negative direction.

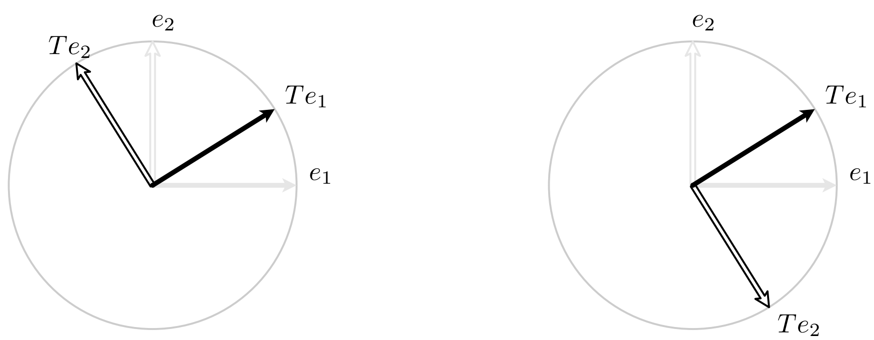

In the first case, we obtain Te_1 and Te_2 from e_1 and e_2 by a rotation. In the second case, something more complicated is going on. Here, as we move a vector u from e_1 towards e_2 and all the way around again to e_1, Tu moves along the arc from Te_1 to Te_2 all the way around again to Te_1, and *in the opposite direction.* Now if we start with two points anywhere on the unit circle and move them around in opposite directions, sooner or later they will meet. At that point we have $Tu = u$. Since T fixes u it fixes the line through u, hence takes points on the line through the origin perpendicular to it into itself. It cannot fix the points on that line, so it must negate them. In other words, T amounts to reflection in the line through u.

Exercise 4.2. *Explain why we can take u to be either of the points half way between e_1 and Te_1.*

Exercise 4.3. *Find a formula for the reflection of v in the line through u.*

5. Rigid motions in 3D

There is one natural way to construct rigid linear motions in 3D. Choose an axis, and choose on it a direction. Equivalently choose a unit vector u, and the axis to be the line through u, with direction that of u.

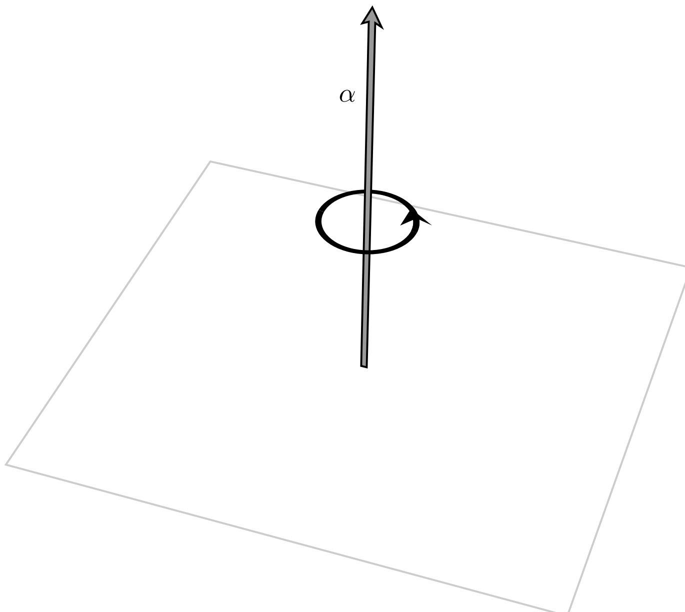

Choose an angle θ. Rotate things around the axis through angle θ, in the positive direction as seen along the axis from the positive direction. This is called an **axial rotation**.

- *The only orientation-preserving linear rigid transformations in 3D are axial rotations.*

We shall give two proofs, one algebraic and the other geometric. But we postpone both until later in the chapter.

6. Calculating the effect of axial rotations

To begin this section, I remark again that to determine an axial rotation we must specify not only an axis but a direction on that axis. This is because the sign of a rotation in $3D$ is only determined if we know whether it is assigned by a left hand or right hand rule. At any rate if choosing a vector along an axis fixes a direction on it. Given a direction on an axis we shall adopt the convention that the direction of positive rotation follows the right hand rule.

So now the question we want to answer is this: *Given a vector $\alpha \neq 0$ and an angle θ. If u is any vector in space and we rotate u around the axis through α by θ, what new point v do we get?* This is one of the main calculations we will make to draw moved or moving objects in $3D$.

There are some cases which are simple. If u lies on the axis, it is fixed by the rotation. If it lies on the plane perpendicular to α it is rotated by θ in that plane (with the direction of positive rotation determined by the right hand rule).

If u is an arbitrary vector, we express it as a sum of two vectors, one along the axis and one perpendicular to it, and then use linearity to find the effect of the rotation on it.

To be precise, let R be the rotation we are considering. Given u we can find its **projection** onto the axis along α to be

$$u_0 = \left(\frac{\alpha \bullet u}{\alpha \bullet \alpha} \right) \alpha$$

Its projection $u_\perp$ is then $u - u_0$. We write

$$\begin{aligned} u &= u_0 + u_\perp \\ Ru &= Ru_0 + Ru_\perp \\ &= u_0 + Ru_\perp \ . \end{aligned}$$

How can we find $Ru_\perp$?

Normalize α so $\|\alpha\| = 1$, in effect replacing α by $\alpha/\|\alpha\|$. This normalized vector has the same direction and axis as α. The vector $u_* = \alpha \times u_\perp$ will then be perpendicular to both α and to $u_\perp$ and will have the same length as $u_\perp$. The plane perpendicular to α is spanned by $u_\perp$ and u_*, which are perpendicular to each other and have the same length. The following picture shows what we are looking at from on top of α.

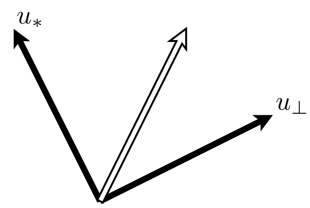

It shows:

- *The rotation by θ takes $u_\perp$ to*

$$Ru_\perp = (\cos\theta)\, u_\perp + (\sin\theta)\, u_* \ .$$

In summary:

(1) Normalize α, replacing α by $\alpha/\|\alpha\|$.

(2) Calculate

$$u_0 = \left(\frac{\alpha \bullet u}{\alpha \bullet \alpha}\right) \alpha \ .$$

(3) Calculate

$$u_\perp = u - u_0 \ .$$

(4) Calculate

$$u_* = \alpha \times u_\perp \ .$$

(5) Finally set

$$Ru = u_0 + (\cos\theta)\, u_\perp + (\sin\theta)\, u_* \ .$$

Exercise 6.1. *What do we get if we rotate the vector $(1, 0, 0)$ around the axis through $(1, 1, 0)$ by $36°$?*

Exercise 6.2. *Write a PostScript procedure with α and θ as arguments and returns the matrix associated to rotation by θ around α.*

7. Eigenvalues and rotations

In this section we shall see the first proof that all linear, rigid, orientation-preserving transformations are axial rotations. It requires the notions of eigenvalue and eigenvector.

If T is any linear operator, an eigenvector of T is a vector $v \neq 0$ such that Tv is a scalar multiple of v:

$$Tv = cv \ .$$

the number c is called the eigenvalue corresponding to v.

If A is a matrix representing T, then the eigenvalues of T are the roots of the characteristic polynomial

$$\det(A - xI)$$

where x is a variable. For a 3×3 matrix

$$A - xI = \begin{bmatrix} a_{1,1} - x & a_{1,2} & a_{1,3} \\ a_{2,1} & a_{2,2} - x & a_{2,3} \\ a_{3,1} & a_{3,2} & a_{3,3} - x \end{bmatrix}$$

and the characteristic polynomial is a cubic polynomial which starts out

$$-x^3 + \cdots \ .$$

For $x < 0$ and $|x|$ large, this expression is positive, and for $x > 0$ and $|x|$ large it is negative. It must cross the x-axis somewhere, which means that it must have at least one real root. Therefore A has at least one real eigenvalue. In 2D this argument fails—there may be two conjugate complex eigenvalues instead.

Let c be a real eigenvalue of T, v a corresponding eigenvector. Since T is a rigid transformation, $\|Tv\| = \|v\|$, or $\|cv\| = \|v\|$. Since $\|cv\| = |c|\|v\|$ and $\|v\| \neq 0$, $|c| = 1$ and $c = \pm 1$.

If $c = 1$, then we have a vector fixed by T. Since T preserves angles, it takes all vectors in the plane through the origin perpendicular to v into itself. Since T reserves orientation and $Tv = v$, the restriction of T on this plane also preserves orientation. Therefore T rotates vectors in this plane, and must be a rotation around the axis through v.

If $c = -1$, then we have $Tv = -v$. The transformation T still takes the complementary plane into itself. Since T preserves orientation in 3D but reverses orientation on the line through v, T reverses orientation on this plane. But then T must be a reflection on this plane. We can find u such that $Tu = u$, and w perpendicular to u and v such that $Tw = -w$. In this case, T is rotation through $180°$ around the axis through u.

8. Finding the axis and angle

If we are given a matrix R which we calculate to be orthogonal and with determinant 1, how do we find its axis and rotation angle? (1) *How do we find its axis?* If e_i is the i-th standard basis vector (one of **i**, **j**, or **k**) the i-th column of R is Re_i. Now for any vector u the difference $Ru - u$ is perpendicular to the rotation axis. Therefore we can find the axis by calculating a cross product $(Re_i - e_i) \times (Re_j - e_j)$ for one of the three possible distinct pairs from the set of indices 1, 2, 3–unless it happens that this cross-product vanishes. Usually all three of these cross products will be non-zero vectors on the rotation axis, but in exceptional circumstances it can happen that one or more will vanish. It can even happen that all three vanish! But this only when A is the identity matrix, in which case we are dealing with the trivial rotation, whose axis isn't well defined anyway.

At any rate, any of the three which is not zero will tell us what the axis is.

(2) *How do we find the rotation angle?*

As a result of part (1), we have a vector α on the rotation axis. Normalize it to have length 1. Choose one of the e_i so that α is not a multiple of e_i. Let $u = e_i$. Then Ru is the i-th column of R.

Find the projection u_0 of u along α, set $u_\perp = u - u_0$. Calculate $Ru_\perp = Ru - u_0$. Next calculate

$$u_* = \alpha \times u_\perp \ .$$

and let θ be the angle between $u_\perp$ and $Ru_\perp$. The rotation angle is θ if $Ru_\perp \bullet u_* >= 0$ otherwise $-\theta$.

Exercise 8.1. *If*

$$R = \begin{bmatrix} 0.899318 & -0.425548 & 0.100682 \\ 0.425548 & 0.798635 & -0.425548 \\ 0.100682 & 0.425548 & 0.899318 \end{bmatrix}$$

find the axis and angle.

9. Euler's Theorem

The fact that every orthogonal matrix with determinant 1 is an axial rotation may seem quite reasonable, after some thought about what else might such a linear transformation be, but I claim that it is not quite intuitive. To demonstrate this, let me point out that it implies that *the combination of two rotations around distinct axes is again a rotation.* This is not at all obvious, and in particular it is difficult to see what the axis of the combination should be. This axis was constructed geometrically by Euler.

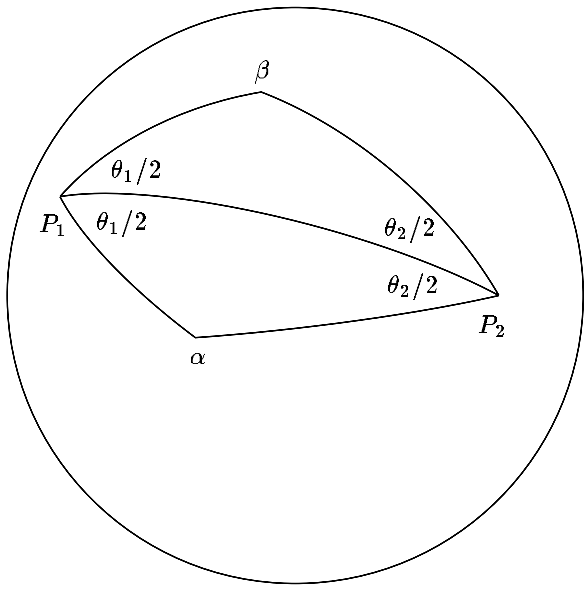

Let P_1 and P_2 be points on the unit sphere. Suppose P_1 is on the axis of a rotation of angle θ_1, P_2 that of a rotation of angle θ_2. Draw the spherical arc from P_1 to P_2. On either side of this arc, at P_1 draw arcs making an angle of $\theta_1/2$ and at P_2 draw arcs making an angle of $\theta_2/2$. Let these side arcs intersect at α and β on the unit sphere. The the rotation R_1 around P_1 rotates α to β, and the rotation R_2 around P_2 moves β back to α. Therefore α is fixed by the composition R_2R_1, and must be on its axis.

Exercise 9.1. *What is the axis of R_1R_2? Prove geometrically that generally*

$$R_1R_2 \neq R_2R_1 \ .$$

Given Euler's Theorem, we can finish our second proof that all linear rigid transformations in 3D are axial rotations by showing the following to be true:

- Any linear rigid transformation can be expressed as the composition of axial rotations.

Let T be the given linear rigid transformation. Let N be the 'north pole' $(0,0,1)$, S the 'south pole' $(0,0,-1)$, E the corresponding 'equator' on the unit sphere. If $TN = N$, then T itself must be just a rotation around the NS axis. If $TN \neq N$, then a single rotation A around an axis through opposite points of E will bring TN up to N, but then TA must be an axial rotation B, and $T = BA^{-1}$.

10. Linear transformations and matrices

There is one point we have been a bit careless about. Suppose T to be a rigid linear transformation. It has been asserted that if $Tv = v$ then T preserves orientation on the plane through the origin perpendicular to v, and if

$Tv = -v$ then it reverses orientation on that plane. Why exactly is this? It depends on a fundamental but difficult point about the relationship between linear transformations and matrices.

There is no canonical way to associate a matrix to a linear transformation, A linear transformation is a geometrical thing—it rotates or reflects or scales or shears in some way. A matrix is in some sense the set of coordinates of the transformation. Something similar happens with vectors which are also geometrical things, possessing, as you are told in physics courses, direction and magnitude. In order to assign coordinates to a vector we must choose first a coordinate system. There is really no best way to do this, and in some sense large classes of coordinate systems are equivalent. Likewise, to assign a matrix to a linear transformation we must first choose a coordinate system. If we do that, say in space, then we can define three special vectors $e_1 = (1, 0, 0)$, $e_2 = (0, 1, 0)$, $e_3 = (0, 0, 1)$. The coordinate system is in turn determined completely by these three vectors, which are called the **basis** of vectors determined by the coordinate system. If v is any vector then it may be expressed as a linear combination $c_1e_1 + c_2e_2 + c_3e_3$. The coordinates of the head of v are the coefficients c_i.

At any rate, we get a matrix from a linear transformation T by setting the entries in its i-th column to be the coordinates of the point Te_i.

Suppose we choose two different coordinate systems, with special vectors $e_\bullet$ and also $f_\bullet$. We get from the first a matrix A associated to T, and from the second a matrix B. What is the relationship between the matrices A and B? We shall not see here a proof, but the result is relatively simple. Let F be the 3×3 matrix whose columns are the coordinates of the vectors $f_\bullet$ in terms of the vectors $e_\bullet$. Then

$$AF = FB \, .$$

The important consequence of this for us now is that the determinant of a linear transformation, which is defined in terms of a matrix associated to it, is independent of the coordinate system which gives rise to the matrix. That is because

$$A = FBF^{-1}, \qquad \det(A) = \det(FBF^{-1}) = \det(F)\det(B)\det(F)^{-1} = \det(B) \, .$$

For our immediate purposes we apply this in this fashion: suppose $Tv = -v$. We choose a set of basis vectors with the first equal to v and the others in the plane perpendicular to v. Since T preserves this plane, we can associate to it a 2×2 matrix A. The 3D matrix of T is then

$$\begin{bmatrix} -1 & 0 & 0 \\ 0 & a & b \\ 0 & c & d \end{bmatrix}, \qquad A = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \, .$$

The determinant of this matrix is then equal to $-\det(A)$. Since T preserves orientation, this must be positive, which implies that $\det(A) < 0$, and T acts upon the perpendicular plane so as to reverse orientation.